

Quantum Generative Adversarial Networks for Learning and Generating Noisy Entangled Quantum States

Albert A. Yelken

Fusion Global Academy, 72 Monroe Center St NW, Suite B, Grand Rapids, MI, 49503, United States

ABSTRACT

Quantum computing hardware is prone to decoherence and control errors, which can reduce entanglement and lead to mixed states. Practical quantum algorithms must extract insights from this noisy data. We test the hypothesis that a compact, circuit-based Quantum Generative Adversarial Network (QGAN) trained against a moving, noise-scheduled target can learn the action of a depolarizing channel and generate a family of noisy entangled two-qubit states that match state-level similarity and reproduce entanglement trends. We implement a compact QGAN in PennyLane with TensorFlow and train it against a moving target formed by Bell-pair states passed through a depolarizing channel whose strength is scheduled during training. The generator and discriminator are variational circuits; learning is assessed by state-level similarity (fidelity and trace distance) and by entanglement measures (concurrence and negativity). Our model learns the action of the noise channel and reproduces a family of noisy, entangled two-qubit states. Our results indicate that QGANs can capture hardware-relevant noise while preserving essential structure in the data. Such models can serve as error-aware state preparers, compact surrogates for device noise, and practical tools for quantum data augmentation and benchmarking on near-term quantum computing hardware. To our knowledge, this is the first systematic QGAN evaluation combining dynamic depolarizing noise schedule, finite-shot/readout-noisy training, and entanglement-aware metrics.

Keywords: Quantum Computing; Quantum Software; Machine Learning; Quantum Noise; Variational Circuits

INTRODUCTION

Quantum computing leverages quantum-mechanical principles, including superposition, entanglement, and interference, to perform information processing

in fundamentally novel ways (1). Entangled states are foundational to quantum protocols, underpinning processes such as quantum teleportation, superdense coding, and critical tests of nonlocality (1, 2). Quantum hardware is vulnerable to decoherence, control inaccuracies, and environmental interactions, which compromise entanglement and lead to mixed states. Practical quantum algorithms must address and extract insights from this noisy data (3, 4).

Quantum Machine Learning (QML) seeks to bridge the gap between classical algorithms and quantum architecture, potentially enhancing expressivity,

Corresponding author: Albert A. Yelken, E-mail: ayel36461@fusionacademy.me

Copyright: © 2025 Albert A. Yelken. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Accepted October 21, 2025

<https://doi.org/10.70251/HYJR2348.3510261032>

sampling efficiency, and computational throughput (3,4). Early QML proposals included quantum versions of support vector machines and principal component analysis. Recently, the focus has shifted to variational circuits for classification and generative modeling. With advancements in quantum processors, QML is poised to leverage the capabilities of noisy intermediate-scale quantum (NISQ) devices, presenting a compelling frontier for enhancing machine learning methodologies (5). Generative Adversarial Networks (GANs) are a deep learning architecture introduced by Goodfellow et al. in 2014. Quantum Generative Adversarial Networks (QGANs) enhance the traditional GANs framework to handle quantum data (7, 8). A generative model that combines coherent structure with incoherent noise has two key benefits: it aids in developing quantum algorithms by simulating realistic hardware and enhances error mitigation and augmentation in quantum machine learning workflows (7, 8).

In this work, we test the hypothesis that a compact, circuit-based QGAN trained against a moving, noise-scheduled target can learn the action of a depolarizing channel and generate a family of noisy entangled two-qubit states that match state-level similarity (fidelity, trace distance) and reproduce entanglement trends. Rather than overfitting an ideal Bell state, the model internalizes the noise channel and yields mixed states reflecting channel-induced degradation. Concretely, we use a depolarizing channel with strength p and a moving-target schedule in which the target state $\rho_{\text{target}}(p)$ is updated by gradually increasing p during training. We address three questions: Can a compact QGAN reproduce target-state similarity across noise strengths p ? Does the learned model track entanglement decay as p increases (via concurrence and negativity)? How do the moving-target schedule and realistic training-time circuit noise affect convergence and final agreement with the targets?

While many previous QGAN works acknowledge noise qualitatively or benchmark under a fixed noise level, fewer provide systematic, curriculum-style sweeps of noise strength with entanglement-aware evaluation (7-10). Our study differs in three practical ways. First, we target noisy entangled two-qubit states explicitly and ramp the depolarizing strength $p \in [0.05, 0.45]$ during training, turning the problem into a moving-target generative task rather than fitting a single fixed distribution. Second, we train under finite shots and readout noise using a three-qubit discriminator (readout ancilla) and a two-qubit generator, mirroring realistic NISQ constraints. Third,

beyond fidelity and trace distance, we track entanglement (as measured by concurrence and negativity) of both the target and generated states, showing that the QGAN not only matches density matrices with a fidelity of ≥ 0.99 at high noise but also reproduces the physical decay and eventual loss of entanglement as noise increases. This emphasis on noise-scheduled training with entanglement-level validation complements prior QGAN demonstrations. It underscores the suitability of adversarial learning for error-aware state preparation and quantum data augmentation on near-term devices. To our knowledge, this is the first systematic QGAN evaluation on noisy entangled states that combines a dynamic depolarizing noise schedule, finite-shot/readout-noisy training, and entanglement-aware metrics.

LITERATURE REVIEW

Quantum Generative Adversarial Networks (QGANs)

QGANs employ two neural networks, a generator and a discriminator, both are implemented as parameterized quantum circuits (7, 8). The generator creates quantum states that aim to match a target quantum state distribution, while the discriminator evaluates these generated states against genuine quantum samples to produce a classification result. The learning process is supported by a quantum-classical feedback loop, which adjusts the circuit parameters, enabling QGANs to effectively approximate complex quantum probability distributions (11, 12). This capability includes modeling entangled multi-qubit states and outputs from noisy quantum channels, tasks that are particularly challenging for classical modeling techniques (9, 10).

Early QGAN Implementations

Several formulations of quantum generative adversarial learning have been proposed in parallel, adapting the adversarial game to parameterized quantum circuits for the generator and, in some cases, the discriminator (7, 8). Early demonstrations focused on learning simple quantum state families (pure states, mixtures, or thermal-like ensembles) from measurement data and on low-dimensional discrete distributions, such as Bars-and-Stripes (BAS) patterns encoded in few-qubit registers. Subsequent works explored hybrid QGANs with a quantum generator and a classical discriminator (or vice versa), as well as fully quantum QGANs in which both players are variational circuits measured in different bases (7-10). Beyond toy distributions, QGAN-style training has been applied to

tasks such as state compression/autoencoding, quantum channel learning, and data-driven state preparation for small molecules or spin models, typically on simulators or few-qubit devices.

Broader Quantum Generative Models

QGANs sit alongside broader quantum generative models, including Quantum Circuit Born Machines and related maximum-likelihood or score-matching methods, which fit a parameterized wavefunction to sample data, as well as Boltzmann-type quantum models trained via contrastive objectives. For state preparation and representation learning, variational algorithms (e.g., VQE-style ansätze), quantum autoencoders, and learning from classical/quantum shadows provide alternative routes that do not rely on an explicit adversarial game. These approaches often optimize fidelity-like objectives or reconstruct reduced statistics, rather than training against an adaptive discriminator (7-10).

Noise and Variational Quantum Algorithms

Because near-term hardware is noisy, prior studies frequently incorporate simple stochastic channels, most commonly depolarizing, amplitude-damping, and phase-damping, either in simulations or as part of hardware runs, and report how noise affects the convergence and expressivity of variational circuits. The literature also documents SPAM (state preparation and measurement) errors, finite-shot effects, and training pathologies such as barren plateaus, which motivate techniques like measurement-error mitigation, zero-noise extrapolation, and probabilistic error cancellation (1).

METHODS AND MATERIALS

We implement a QGAN in PennyLane with TensorFlow to learn and generate noisy entangled two-qubit states (11, 13). We analyze convergence behavior using via fidelity F and trace distance T , explore the interplay between adversarial dynamics and noise, and quantify entanglement decay through concurrence C and negativity N (2, 14).

Target Distribution and Noise Model

We consider adversarial learning of noisy entangled two-qubit states. The canonical Bell state

$|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ represents the simplest maximally

entangled two-qubit system (1). The target distribution is

a Bell state $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$, $\rho_{\text{Bell}} = |\Phi^+\rangle\langle\Phi^+|$, subjected

independently on each qubit to a single-qubit

depolarizing channel $\mathcal{D}_p(\rho) = (1-p)\rho + \frac{p}{2}I$. At training

epoch e , we set the depolarizing parameter to follow a linear schedule $p_e \in [0.05, 0.45]$ yielding the noisy target state $\rho_{\text{real}}(p_e) = (\mathcal{D}_{p_e} \otimes \mathcal{D}_{p_e})(\rho_{\text{Bell}})$. This ensures that the QGAN must adapt to a moving target distribution as noise increases over training.

QGAN Architecture

The QGAN comprises two parameterized quantum circuits: a generator G_θ and a discriminator D_ϕ . The discriminator receives either a noisy Bell sample from $\rho_{\text{real}}(p_e)$ or a generated sample from G_θ , and outputs an acceptance probability. Training alternates between optimizing ϕ and θ via gradient updates.

Generator Ansatz

The generator acts on two qubits and uses a 9-parameter circuit. It begins with a Hadamard on qubit 0, then applies local rotations on both qubits $R_X(\theta_0) R_Y(\theta_1) R_Z(\theta_2) \otimes R_X(\theta_3) R_Y(\theta_4) R_Z(\theta_5)$ followed by an entangling gate $\text{CNOT}_{0 \rightarrow 1}$ and a final local rotation layer on qubit 0 $R_X(\theta_6) R_Y(\theta_7) R_Z(\theta_8)$. The resulting noiseless generator state is $\rho_G^{\text{pure}}(\theta) = U_G(\theta) |00\rangle\langle 00| U_G^\dagger(\theta)$ and, to match the training distribution, the output is exposed to the same depolarizing noise at the current schedule value p_e , yielding $\rho_G(\theta, p_e) = (\mathcal{D}_{p_e} \otimes \mathcal{D}_{p_e})(\rho_G^{\text{pure}}(\theta))$.

Discriminator Ansatz

The discriminator uses three qubits, two data qubits and one ancilla, and is built as a single parametric block. It begins with local rotations on each qubit, $\text{Rot}(\phi_{0,2})$ on qubit 0, $\text{Rot}(\phi_{3,5})$ on qubit 1, $\text{Rot}(\phi_{6,8})$ on the ancilla. These are followed by two entangling operations that couple the data qubits to the ancilla, $\text{CNOT}_{0 \rightarrow 2}$ and $\text{CNOT}_{1 \rightarrow 2}$, and a final ancilla rotation $\text{Rot}(\phi_{9,11})$. The circuit outputs the expectation value of the ancilla in the basis, which we convert to an acceptance probability, $p_{\text{acc}} = \frac{1}{2}(\langle Z_{\text{anc}} \rangle + 1)$. To model realistic measurement imperfections, we include a readout bit-flip channel on the ancilla with probability $q = 0.02$.

Cost Functions

Let $p_r(\phi)$ denote the discriminator's acceptance

probability for real data, and $p_f(\theta, \phi)$ for generated data. We adopt the standard cross-entropy losses:

$$\mathcal{L}_D(\phi) = -[\log p_r(\phi) + \log(1 - p_f(\theta, \phi))], \mathcal{L}_G(\theta) = -\log p_f(\theta, \phi)$$

Training Procedure

Training was carried out for a total of 200 epochs. In each epoch, the update schedule consisted of $K_D = 3$ discriminator steps followed by $K_G = 1$ generator steps. Parameters were optimized using the Adam optimizer with a learning rate 0.05, $\beta_1 = 0.5$, and $\beta_2 = 0.9$ (15). Circuits were executed on a mixed-state simulator with 2000 shots, and gradients were computed using the parameter-shift rule. The noise applied during training followed the linearly increasing depolarizing schedule $p \in [0.05, 0.45]$, ensuring that the QGAN adapted continuously to progressively noisier target states.

Evaluation Metrics

We evaluate performance against both the noisy target state $\rho_{\text{real}}(p_e)$ and the ideal Bell reference ρ_{Bell} . As a primary similarity measure we use the Uhlmann

fidelity, $F(\rho, \sigma) = \left(\text{Tr} \sqrt{\sqrt{\rho} \sigma \sqrt{\rho}} \right)^2$, which lies in $[0, 1]$

and equals 1 only for identical states; complementary to this we report the trace distance, $T(\rho, \sigma) = \frac{1}{2} \|\rho - \sigma\|_1$, which quantifies optimal single-shot discrimination advantage and also lies in $[0, 1]$. Because our targets are entangled and the noise schedule destroys entanglement as p_e increases, we additionally track two-qubit entanglement monotones: the concurrence, $C = \max(0, \lambda_1 - \lambda_2 - \lambda_3 - \lambda_4)$ where $\{\lambda_i\}$ are the square roots (sorted nonincreasingly) of the eigenvalues of

$$\rho(\sigma_y \otimes \sigma_y) \rho^*(\sigma_y \otimes \sigma_y), \text{ and the negativity, } \mathcal{N} = \sum_{\mu_i < 0} |\mu_i|,$$

where $\{\mu_i\}$ are the eigenvalues of the partial transpose of ρ . For a Bell state, one has $\mathcal{C}(\rho_G)$ and $\mathcal{N}(\rho_G)$ to those of $\rho_{\text{real}}(p_e)$ provides a check that the model reproduces not only density-level similarity but also the correct entanglement structure. In all plots, we report epoch-wise trajectories of F and T to $\rho_{\text{real}}(p_e)$ and to ρ_{Bell} , together with concurrence and negativity for both real and generated states, and we visualize representative density matrices (real and imaginary parts) to qualitatively corroborate these quantitative trends (1, 16). For clarity, metrics are computed from simulator density matrices (no shots) to remove estimator variance, while training itself is performed with finite shots and readout noise; thus, the reported metrics

reflect intrinsic state closeness under the imposed noise schedule rather than measurement fluctuations.

RESULTS

Training dynamics and stability

The adversarial game remained well-behaved across the entire noise schedule $p \in [0.05, 0.45]$. The discriminator and generator losses fluctuated around steady plateaus (typical ranges ~ 1.30 – 1.46 for D and ~ 0.60 – 0.75 for G), with no runaway gradients or collapse. The discriminator’s acceptance probabilities for real and generated states both hovered near 0.5 throughout training, e.g., at the final epoch $P(\text{real} \rightarrow \text{real}) \approx 0.52$ and $P(\text{fake} \rightarrow \text{real}) \approx 0.50$. Small, rapid oscillations were observed (2000 shots; readout bit-flip $q \approx 0.02$) (Figure 1).

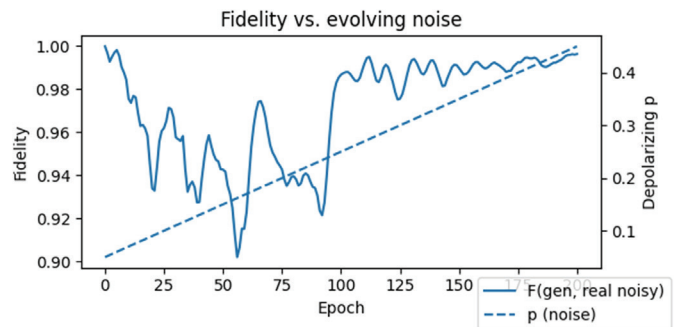


Figure 1. Fidelity vs. evolving noise (tracking a moving noisy target). Uhlmann fidelity $F(\rho_G, \rho_{\text{real}}(p))$ between the generator output and the noise-scheduled target across epochs. The dashed curve shows the depolarizing parameter p (right axis) ramped linearly during training. Fidelity dips as the model adapts to the moving target, then rises and stabilizes near 1, indicating close tracking of the evolving mixed-state distribution.

Similarity to the noise-scheduled target (Fidelity and Trace Distance)

We evaluate against the noisy target $\rho_{\text{real}}(p)$ and the ideal Bell state $\rho_{\text{real}}(p_e)$ using fidelity $F(\rho, \sigma)$ and trace distance $T(\rho, \sigma)$. As p increases, the QGAN tracks the target closely and improves in both metrics: by the end of training, we obtain $F(\rho_G, \rho_{\text{real}}) = 0.9963$ and $T(\rho_G, \rho_{\text{real}}) = 0.0475$ at $p = 0.45$. Along the schedule, representative checkpoints were: epoch 20 ($p = 0.09$) with $F = 0.9340$, $T = 0.2418$; epoch 100 ($p = 0.25$) with $F = 0.9867$, $T = 0.0981$; and an epoch 180 ($p = 0.41$) with $F = 0.9942$, $T =$

0.0604. ($F(\rho_G, \rho_{\text{real}})$) decreased to (≈ 0.36) at ($p = 0.45$), while ($T(\rho_G, \rho_{\text{real}})$) increased to (≈ 0.64). At ($p = 0.05$), the initialization yielded $F \approx 0.9999$ (Figure 2).

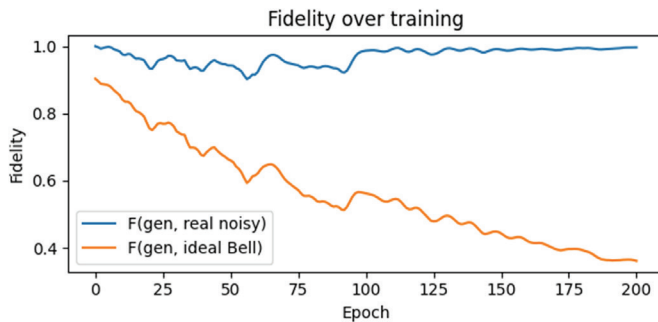


Figure 2. Fidelity to the noisy target versus the ideal Bell state. Fidelity of the generated state to the noise-scheduled target (blue) remains high and improves late in training, whereas fidelity to the ideal Bell state (orange) decreases steadily. This divergence confirms that the model learns the channel-induced mixed-state family rather than overfitting a pure Bell state.

Early-epoch alignment (Initialization effect)

We monitor entanglement with concurrence and negativity. Both the target and the generated states exhibit monotone entanglement decay with p , vanishing near the high-noise end of the schedule. By the final epoch, we observe $C(\rho_{\text{real}}) \approx 0$, $N(\rho_{\text{real}}) \approx 0$, and the generator matches this trend, $C(\rho_G) \approx 0$, $N(\rho_G) \approx 0$ (Figure 3).

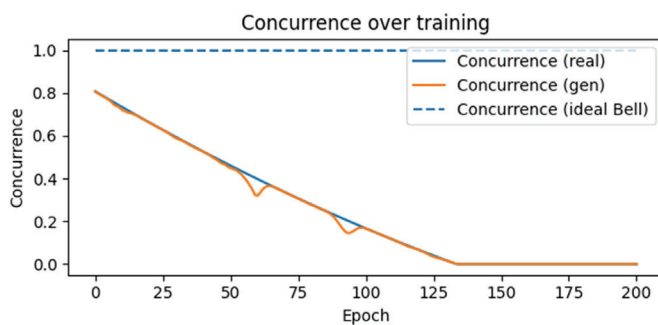


Figure 3. Concurrence over training (entanglement decay). Concurrence $C(\rho)$ for the target and generated states under the increasing-noise schedule. The two curves closely coincide and decrease to $C \approx 0$ at late epochs; a dashed line at $C \approx 1$ marks the ideal Bell reference. The behavior matches the expected depolarization-driven loss of two-qubit entanglement.

Entanglement trends (Concurrence and Negativity)

The largest eigenvalue of the noisy target provides a pure-state fidelity ceiling; for our two-qubit depolarizing-on-each-qubit model at $p = 0.45$, the bound is $\lambda_{\text{max}} \approx 0.37$. The final model achieved final $F \approx 0.9963$.

DISCUSSION

In this work, we develop and evaluate a circuit-based QGAN for learning and generating noisy entangled qubit pairs under a depolarizing channel. Training dynamics indicate a near-equilibrium adversarial game. Acceptance probabilities for real and generated inputs remained close to 0.5, which is consistent with a Nash-like balance in which the discriminator cannot reliably distinguish the two distributions. The small oscillations in losses and acceptance traces are consistent with finite-shot sampling ($N = 2000$) and a readout bit-flip channel ($q \approx 0.02$), rather than genuine instability of the optimization.

The fidelity and trace-distance trends show that the generator learns the noise-scheduled target, not a clean Bell state. As p increases, $F(\rho_G, \rho_{\text{real}}(p))$ improves while $F(\rho_G, \rho_{\text{Bell}})$ decreases and $T(\rho_G, \rho_{\text{Bell}})$ increases, reflecting deliberate divergence from the ideal state as the target becomes noisier. The very high fidelity observed at initialization for the smallest noise level $F \approx 0.9999$ at $p = 0.05$ can arise when the random seed and early schedule point place the ansatz near a good local match for a weakly mixed state; subsequent training steps then follow the schedule as p grows.

Entanglement diagnostics (concurrence and negativity) decay monotonically with p for both target and generated states and approach zero at high noise. This agreement implies that the model reproduces not only the classical mixture statistics but also the loss of quantum correlations induced by the channel.

The pure-state fidelity ceiling provides a useful consistency check. At $p = 0.45$, the target's largest eigenvalue gives a ceiling $\lambda_{\text{max}} \approx 0.37$ for any pure-state approximation, yet the trained model attains $F \approx 0.9963$. This is expected because the generator output is passed through the same noise channel before evaluation, producing mixed states. The result, therefore, exceeds the pure-state ceiling without contradiction and underscores that `\text{explicit channel modeling is essential}` for accurate matching.

Why does $F(\rho_G, \rho_{\text{real}})$ improve with p ? As depolarizing strength increases, the target distribution approaches an isotropic mixture, which is simpler to

approximate with a shallow circuit; this naturally yields higher fidelities later in the schedule.

A more quantitative explanation of the rising fidelity with increasing depolarization is as follows. Independent single-qubit depolarizing channels $\mathcal{D}_p^{\otimes 2}$ acting on $|\Phi^+\rangle$ produce a Bell-diagonal (isotropic)

$$\text{target } \rho_{\text{real}}(p) = a(p) |\Phi^+\rangle\langle\Phi^+| + \frac{1-a(p)}{3} \sum_{\beta \in \{\Phi^-, \Psi^+, \Psi^-\}} |\beta\rangle\langle\beta|,$$

$$a(p) = 1 - 2p + \frac{4}{3}p^2$$

Thus, the spectrum is $\{a(p), \frac{1-a(p)}{3}, \frac{1-a(p)}{3}, \frac{1-a(p)}{3}\}$; in particular, the largest eigenvalue is $a(p)$ (e.g., $a(0.45) \approx 0.37$, which sets the “pure-state ceiling” used above.

Crucially, the generator output is also passed through the same $\mathcal{D}_p^{\otimes 2}$ before comparison. Increasing is equivalent to composing an additional depolarizing map $\mathcal{D}_{p+\Delta} = \mathcal{D}_\Delta \circ \mathcal{D}_p$, and Uhlmann fidelity is monotone under CPTP maps; therefore, even for a fixed pre-noise generator state, $F(\rho_G(p+\Delta), \rho_{\text{real}}(p+\Delta)) \geq F(\rho_G(p), \rho_{\text{real}}(p))$

This guarantees systematic upward pressure on F as the shared depolarization increases, independent of learning success.

An equivalent view in the Pauli basis clarifies the rate of this effect. Depolarization shrinks every

single-qubit Bloch component by $r(p) = 1 - \frac{4p}{3}$; two-qubit correlation terms are therefore suppressed by $r(p)^2$.

Writing $\rho_G(p) = \frac{\mathbb{I}}{4} + r(p)^2 \Delta_G$ and $\rho_{\text{real}}(p) = \frac{\mathbb{I}}{4} + r(p)^2 \Delta_{\text{Bell}}$, the mismatch becomes $\rho_G(p) - \rho_{\text{real}}(p) = r(p)^2 (\Delta_G - \Delta_{\text{Bell}})$

Hence, the trace distance contracts approximately as $T(p) \propto r(p)^2$, and by the Fuchs–van de Graaf bound $1 - \sqrt{F} \leq T$, the fidelity necessarily increases toward unity as p grows. Practically, the family $\rho_{\text{real}}(p)$ is fully described by the single weight $a(p)$, so the learning problem becomes progressively lower-dimensional and easier for a shallow generator. This explains why $F(\rho_G, \rho_{\text{real}})$ rises later in the schedule, even while $F(\rho_G, \rho_{\text{Bell}})$ and entanglement diagnostics decline: higher p drives both states toward the common identity component, inflating fidelity without recovering lost quantum correlations.

The intentionally minimal architecture (two data qubits, one ancilla, shallow entangling structure) promotes smooth gradients and stable training but limits expressivity for structured or higher-rank noise

models. While shot noise and readout errors perturb the adversarial traces, the density-matrix evaluation of F, T, C, N mitigates measurement noise in the reported metrics. Future extensions should test richer channels, deeper ansätze, and cross-device generalization.

These findings demonstrate that QGANs can function as error-aware state preparers and quantum data-augmentation models for near-term hardware. The models excel at simulating device-induced decoherence rather than achieving pure entangled-state recovery.

CONCLUSION

We have shown, under realistic finite-shot and readout noise, that a modest QGAN can faithfully learn a parameterized family of noisy entangled states, reaching a final fidelity of 0.9963 and a trace distance of 0.0475 at high depolarization. Beyond matching a target, these results indicate that QGANs are effective hardware-aware error-mitigation tools: by learning the device-induced channel and generating states that reflect it, the model reduces sim-to-real mismatch and enables bias-aware estimation or pre-compensation in downstream experiments. In this role, the QGAN acts as an error-aware state preparer—producing mixed states with the correct coherent structure and incoherent noise—and as a quantum data-augmentation engine that reproduces the device’s statistics for training and validation. Because the generator learns distributions, not just single states, it also provides a practical mechanism for loading generic probability distributions into quantum states, supplying informative priors for variational algorithms and QML models. Taken together, our results suggest that QGANs bridge the gap between idealized theory and noisy hardware, strengthening QML pipelines that must operate with realistic, device-specific noise.

Next steps should expand scaling and expressivity (deeper ansätze, richer discriminator measurements, alternative losses such as Wasserstein/MMD) and richer noise models (correlated multi-qubit errors, amplitude damping, non-Markovian and time-varying drift). We see a particular opportunity in deploying the full loop on hardware (superconducting or trapped-ion), benchmarking wall-clock cost, sample complexity, and robustness to calibration drift, and quantifying the benefit of learned error-aware preparation for real algorithms. Application-wise, QGANs are well-suited for building generative priors from quantum data, e.g., quantum-chemistry states, tomography data sets, or

experimental measurement logs, to seed or augment downstream QML tasks (classification, regression, anomaly detection). Finally, learning device channels in situ opens a path to closed-loop error mitigation (channel tracking and pre-compensation) and cross-device generalization, where a single generative model is adapted to multiple backends while preserving essential structure in the data.

FUNDING SOURCES

The author received no external funding for this work.

CONFLICT OF INTERESTS

The author declares that there is no conflict of interest related to this work.

REFERENCES

- Nielsen MA, Chuang IL. Quantum Computation and Quantum Information. 10th Anniversary ed. Cambridge: Cambridge University Press. 2010.
- Wootters WK. Entanglement of formation of an arbitrary state of two qubits. *Phys Rev Lett.* 1998; 80: 2245-8. <https://doi.org/10.1103/PhysRevLett.80.2245>
- Biamonte J, Wittek P, Pancotti N, et al. Quantum machine learning. *Nature.* 2017; 549: 195-202. <https://doi.org/10.1038/nature23474>
- Schuld M, Petruccione F. Supervised Learning with Quantum Computers. *Cham: Springer.* 2018. <https://doi.org/10.1007/978-3-319-96424-9>
- Preskill J. Quantum computing in the NISQ era and beyond. *Quantum.* 2018; 2: 79. <https://doi.org/10.22331/q-2018-08-06-79>
- Goodfellow I, Pouget-Abadie J, Mirza M, et al. Generative adversarial nets. In: Advances in Neural Information Processing Systems (NeurIPS). 2014. arXiv:1406.2661 <https://doi.org/10.48550/arXiv.1406.2661>
- Dallaire-Demers P-L, Killoran N. Quantum generative adversarial networks. *Phys Rev A.* 2018; 98: 012324. <https://doi.org/10.1103/PhysRevA.98.012324>
- Lloyd S, Weedbrook C. Quantum generative adversarial learning. *Phys Rev Lett.* 2018; 121: 040502. <https://doi.org/10.1103/PhysRevLett.121.040502>
- Zoufal C, Lucchi A, Woerner S. Quantum generative adversarial networks for learning and loading random distributions. *npj Quantum Inf.* 2019; 5: 103. <https://doi.org/10.1038/s41534-019-0223-2>
- Hu L, Wu S-H, Cai W, et al. Quantum generative adversarial learning in a superconducting quantum circuit. *Sci Adv.* 2019; 5: eaav2761. <https://doi.org/10.1126/sciadv.aav2761>
- Bergholm V, Izaac J, Schuld M, et al. PennyLane: automatic differentiation of hybrid quantum-classical computations. arXiv preprint arXiv:1811.04968. 2018.
- Schuld M, Bergholm V, Gogolin C, Izaac J, Killoran N. Evaluating analytic gradients on quantum hardware. *Phys Rev A.* 2019; 99: 032331. <https://doi.org/10.1103/PhysRevA.99.032331>
- Abadi M, Barham P, Chen J, et al. TensorFlow: a system for large-scale machine learning. In: Proc OSDI. 2016.
- Vidal G, Werner RF. Computable measure of entanglement. *Phys Rev A.* 2002; 65: 032314. <https://doi.org/10.1103/PhysRevA.65.032314>
- Kingma DP, Ba J. Adam: a method for stochastic optimization. arXiv preprint arXiv:1412.6980. 2014.
- Jozsa R. Fidelity for mixed quantum states. *J Mod Opt.* 1994; 41 (12): 2315-23. <https://doi.org/10.1080/09500349414552171>